

John J. Poklemba
COMSAT Laboratories
Clarksburg, Maryland 20871

INTRODUCTION

In this paper the design of the Programmable Digital Modem (PDM) will be outlined. The PDM will be capable of operating with numerous modulation techniques including: 2-, 4-, 8-, and 16-ary phase shift keying (PSK), minimum shift keying (MSK), and 16-ary quadrature amplitude modulation (QAM), with spectral occupancy from 1.2x to 2x the data symbol rate. It will also be programmable for transmission rates ranging from 2.34 to 300 Mbit/s, where the maximum symbol rate is 75 Msymbol/s. Furthermore, these parameters will be executable in independent burst, dependent burst, or continuous mode. In dependent burst mode the carrier and clock oscillator sources are common from burst to burst.


To achieve as broad a set of requirements as these, it is clear that the essential signal processing must be digital. In addition, to avoid hardware changes when the operational parameters are changed, a fixed interface to an analog intermediate frequency (IF) is necessary for transmission; and, common system level architectures are necessary for the modulator and demodulator. Lastly, to minimize size and power as much of the design as possible will be implemented with application specific integrated circuit (ASIC) chips.

MODULATOR ARCHITECTURE AND DESIGN


Baseband vs IF Digital-to-Analog Sample Conversion

Should the modulator output analog samples at baseband or IF? To answer this, the restrictions caused by the digital-to-analog (D/A) conversion device will first be examined. A D/A converter is inherently a sample-and-hold device that imposes a lowpass $\sin(x)/x$ envelope on the baseband output spectrum and its replicas. This effect is shown in Figure 1a for the integer minimum Nyquist sample rate of two samples/symbol (s/s) and square root 40-percent raised cosine spectral shaping. To support most of the two-dimensional modulation formats listed above, four complex s/s or equivalently two in-phase and two quadrature channel s/s are required. The gap between the main lobe and the first replicated spectra allows a practical analog reconstruction filter to be used, and the D/A stopband notches provide inherent filtering as they occur in the center of the replicated spectra.

To convert the digital baseband samples directly to an IF output at a minimum number of s/s implies that their spectra be shifted up in frequency. To avoid restricting the upper data rate of operation, 3 s/s is the minimum that can be used for IF sampling as shown in Figure 1b. Because of the spectral shift, the D/A converter would cause a considerable amount of amplitude skew across the IF passband; and the first replicated image, centered just above $2R_s$, is very close to the desired lobe, centered just below R_s .



a. Baseband Sample Conversion at 2 Samples/Symbol



b. IF Sample Conversion at 3 Samples/Symbol

Figure 1. D/A Aperture Effects

So even at the minimum bandpass sample rate, it is very difficult to filter out the replicated spectra. Hence, it's clear that for a given speed capability in the digital hardware, baseband sampling will achieve higher data rate operation. Thus, at such high speeds, the most effective way to process the data is with a minimum integer number of samples per symbol with parallel in-phase and quadrature (I and Q) channels at baseband, and analog quadrature carrier mixing for conversion to an IF.

To accommodate multirate operation, the sample rate into the D/A converter will always be within the octave range of 75-150 Msample/s, regardless of the data rate; and the number of samples per symbol will always be a power of two. In this manner, the sample clock replicated spectra of Figure 1a can be removed over the entire symbol rate range of operation with a single analog reconstruction filter. Moreover, the highest symbol rate range is 37.5-75 Msymbol/s at two s/s. The next octave range down is then 18.75-37.5 Msymbol/s at four s/s, and so on.

The replication removal filter must pass as much of the main lobe at the maximum symbol rate ($R_s = 75$ Msymbol/s) as possible, while rejecting the low end of the first replicated lobe at a symbol rate an octave below the maximum ($R_s = 37.5$ Msymbol/s). A good compromise, determined in conjunction with the bit error rate (BER) simulations, is an

*This work was funded under NASA Lewis contract NAS3-25715.

elliptic lowpass filter with a 0.2 dB equiripple passband extending from DC to 48 MHz, with a stopband beginning at 64 MHz of minimum attenuation greater than 30 dB. The sample-and-hold effect of the D/A provides additional filtering to suppress the sample clock replications below 40 dB. To avoid additional analog hardware, group delay dispersion in the replication removal filter will be compensated with digital processing.

A block diagram of the basic modulator architecture is given in Figure 2. The modulator is divided into a digital baseband processor with an analog quadrature carrier IF. The primary function of the baseband processor is to spectrally shape or filter the data in a bandwidth efficient manner, and to convert it to a baseband quadrature format prior to carrier modulation. The quadrature format supports nearly any modulation format that can be represented in a two-dimensional signal space, and the parallel I and Q channels support higher rate operation. The analog portion of the modulator then performs the function of translating the I and Q data representation on to cosine and sine carriers, respectively.

Transmit Spectral Shaping

To achieve the best BER performance possible, it would be desirable to digitally implement and match the transmit and receive filter spectra with a square root Nyquist characteristic, assuming that the remaining filtering functions in the transmission link are transparent. However, in general, the transmit and receive data filters cannot be matched and must be predistorted to account for replication removal, IF, and anti-aliasing filters as well as transmission link impairments.

Because of the strict magnitude and phase constraints for Nyquist data filters, the most appropriate digital filter implementation is the finite impulse response (FIR), which inherently has linear phase. A greatly simplified equivalent implementation is possible because the transmit symbols have relatively few deterministic levels; i.e., BPSK, QPSK,

and MSK only require two input levels. The reduced complexity implementation involves a memory table lookup. A brief description is as follows. Input data symbols are read into a shift register whose length is equal to the number of symbols in the impulse response aperture to be represented. To determine the transmit impulse response, all of the link frequency responses are cascaded, and a discrete Fourier transform (DFT) is employed to compute the predistorted samples. A fast Fourier transform (FFT) is not used because, in general, the sample sets are not a power of N . The symbol patterns in the shift register change every symbol time, so for each symbol pattern there is a unique set of precomputed sample values that will be clocked out of the memory. That is, within a given symbol pattern, there are N unique samples per symbol. The memory size required is determined from

$$M \cdot L \cdot N \quad (1)$$

where

- M = number of in-phase or quadrature symbol amplitude levels required
- L = length of the filtering aperture in symbol times
- N = number of samples per symbol.

Hence, the memory size increases linearly with the number of s/s, but geometrically versus impulse response aperture length and the number of I or Q amplitude levels. For example, a 16-PSK signal constellation will be represented with eight I/Q levels (± 4); whereas QPSK requires only two I/Q levels. Several permutations of the maximum memory sizes required are listed in Table 1 for 32 s/s. The common achievable size for all of the modulation techniques is indicated in parentheses, 131K bytes. Approximate carrier spacings that may be supported are also listed.

The best combination of high density and speed memory currently available is 65K x 4 with an access time of 8 ns, which when setup, hold, and skew times are included, provides a small amount of timing margin for operation at

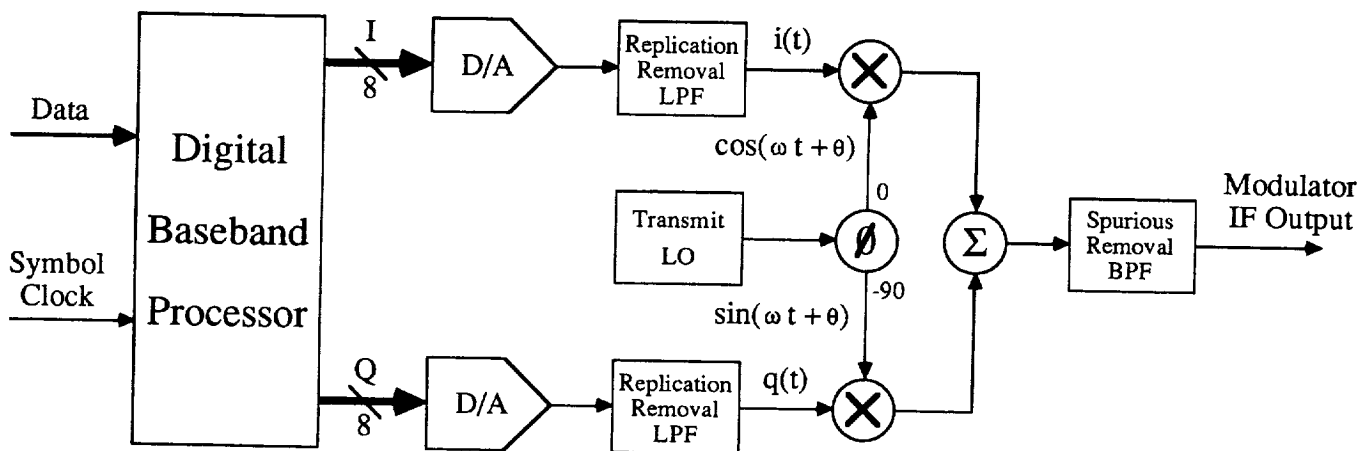


Figure 2. Basic Modulator Architecture

Table 1. Maximum I or Q Channel Memory Requirements at 32 Samples/Symbol

MODULATION TECHNIQUE	NUMBER OF SIGNAL LEVELS	APERTURE LENGTH (SYMBOLS)					
		3	4	5	6	8	12
BPSK, MSK, QPSK	2(± 1)	256	512	1k	2k	8.2k	(131k)
8-PSK, 16-QAM	4($\pm X, \pm Y$)	2k	8k	33k	(131k)	2.1M	
16-PSK	8 ($\pm A, \pm D; \pm B, \pm C$)	16.4k	(131k)	1.0M	8.4M		
Carrier Spacing (R_s Multiples)		1.9	1.8	1.6	1.4	1.3	1.2

75 Msymbol/s (13.3 ns). For 8-bit resolution, four of these chips are required in each of the I and Q channels, along with the 12-symbol shift register. This is considerably simpler than an equivalent 384-tap FIR filter implementation with its incumbent set of digital multiplies and sums. The 8-bit output resolution for the memory results in good spectral quantization noise, which is >40 dB down over the range of rates desired.

DEMODULATOR ARCHITECTURE AND DESIGN

IF vs Baseband Analog-to-Digital Sample Conversion

The issue of sampling directly at IF vs conversion to baseband prior to sampling will now be analyzed separately for the demodulator. With IF sampling, the IF center frequency will scale with the data rate unless a noninteger number of samples per symbol or more complex processing is used. To handle a noninteger number of samples per symbol, an interpolating filter is needed. In the demodulator, the interpolating filter would basically perform two functions. It converts asynchronous samples to synchronous samples at two samples per symbol; such that over each symbol interval, one of the samples occurs at the data detection sample point, while the other is at the average value of the zero crossings for symbol timing recovery. However, an interpolating filter is hardware intensive and speed restrictive. Furthermore, to operate at 75 Msymbol/s suggests that the lowest IF center frequency be at least 75 MHz, or more suitably 140 MHz. A half-cycle of the carrier sinusoid at this rate is about 3.5 ns. The narrowest sampling aperture on currently available analog-to-digital (A/D) converters is on the order of 1.5 to 2 ns. Hence, the width of the sampling aperture is approximately one-half of the slowest practical positive or negative carrier excursion. This imposes a lowpass $\sin(x)/x$ envelope on the incoming bandpass spectra, as was illustrated in Figure 1. For a 1.75-ns aperture, the $\sin(x)/x$ envelope is about 1 dB down at 140 MHz, so sampling at IF would also cause a variable amplitude skew across the passband for the higher operational data rates. As a result of limitations due to the A/D sampling aperture and interpolating filter realizations, the receive bandpass signal will be down converted with

carriers in phase quadrature for subsequent sampling at baseband.

The requirements for the anti-aliasing filter to limit the incoming bandwidth prior to A/D conversion are very comparable to those for the replication removal filter in the modulator. For example, the bulk of the main spectral lobe must be passed at the maximum symbol rate, which extends from DC to 52.5 MHz for a 40-percent Nyquist channel. In addition, the filter must restrict the incoming noise bandwidth to half the minimum sample rate to avoid aliasing at the higher data rates. For this and other reasons which will be explained subsequently, the minimum sample rate on the demodulator, 100 Msample/s is higher than that on the modulator, 75 Msample/s. Previous simulations have shown that 30 dB stopband attenuation is sufficient to have negligible impact on BER, and that greater attenuation merely makes it more difficult to compensate for the filter's delay dispersion. Hence, for simplicity, the anti-aliasing filter will be designed with identical parameters as the modulator replication removal filter. This also allows for a common IF hybrid or MMIC to be developed for use in both the modulator and demodulator.

Demodulator Block Diagram

The basic demodulator structure is given in Figure 3. Note the interdependence of the acquisition estimate processor, the data detection, and the recovery loops. A GaAs ASIC chip is currently being developed that will contain two programmable MACs. It will be capable of being reconfigured to operate in nine separate locations in the demodulator. The ASIC multipliers will be 8 x 8 with a 16-bit barrel shifted output, and the accumulators will be 24 bits with 16-bit preloading. All of the required ASICs will be capable of 150-Msample/s pipeline operation.

Receive Data Detection Filter

The most potentially hardware-intensive function in the demodulator is the receive data detection filter. A memory-based structure is not feasible because of the large number of input quantization levels due to channel impairments and noise. A minimum complexity FIR filter with a reduced or decimated output sample rate is desired. This can be achieved with a very high-speed multiplier-accumulator (MAC), where each accumulator output sample corresponds to a weighted average of a set of incoming samples. Since the output of this filter will feed all of the remaining processing stages necessary in the demodulator, it has been dubbed the "pre-averager" data filter. Separate even and odd MACs are required because the input sample sets that the pre-averager must process are overlapping, as shown in Figure 4 [1]. The even samples are used for data detection, carrier recovery, and gain control; whereas the odd samples provide symbol timing recovery. As indicated, the averages are taken over N samples in one-symbol intervals. So, in effect, the pre-averager impulse response extends over a one-symbol aperture. However, BER simulations with adjacent channels on 1.4x the symbol rate spacings have shown that a one-symbol aperture is not adequate, regardless of the weighting function employed. What is necessary is a

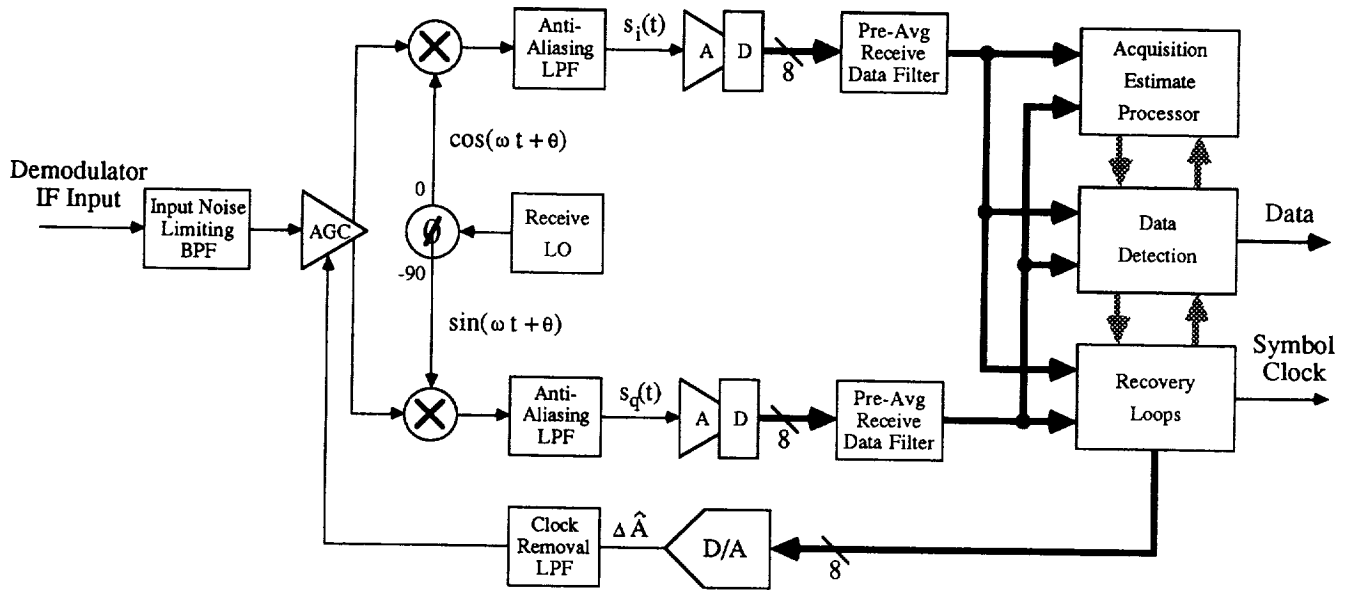


Figure 3. Basic Demodulator Architecture




Figure 4. Overlapping Pre-Averaged Sample Sets

sharper rolloff filtering function that has a stopband in the region above $0.7 R_S$ (half the center-to-center carrier spacing) to remove adjacent channel interference and noise.

Receive Data Filter Impulse Response Derivation

From an implementation point of view, the most straightforward way to modify the poor adjacent channel rejection (ACR) capability of the one-symbol aperture pre-averager is to increase its aperture to two symbols, with 50 percent overlapping averaging intervals. Next, it would be

desirable to find a strictly time-limited two-symbol-long impulse response, with a stopband above $0.7 R_S$. Proceeding to the sampled frequency domain, a very general Nyquist filtering function may be defined to satisfy this condition for two s/s as follows

$$\begin{aligned} H(0) &= 1.0 \\ H(1) &= 0.5 \\ H(2) &= 0.0 \\ H(3) &= 0.5 \end{aligned} \quad (2)$$

where R_S has been normalized to 2.

These four frequency domain samples at two s/s will yield four time domain samples that extend over a two-symbol aperture. Using the definition of the inverse DFT,

$$h(n) = \frac{1}{N} \sum_{k=0}^{N-1} H(k) \exp(j 2\pi kn/N), \quad 0 \leq n \leq N-1 \quad (3a)$$

on the values in equation (2) yields a raised cosine pulse:

$$h(n) = \frac{1}{4} \{ 1 + 0.5 [\exp(j\pi n/2) + \exp(j3\pi n/2)] \} \quad (3b)$$

$$= \frac{1}{2} \left[\frac{1 + \cos(\pi n/2) \exp(-j\pi n)}{2} \right] \quad (3c)$$

$$= \frac{1}{2} \left[\frac{1 + \cos(\pi n/2)}{2} \right] \quad (3d)$$

where the exponential phase term is dropped from the last equality because the cosine term is zero for n -odd, and it has no effect for n -even.

Extensive BER simulations have shown the raised cosine pulse (RCP) impulse response of equation (3d) to be substantially more effective than truncated square root Nyquist impulse responses in providing good adjacent channel rejection, for a two-symbol aperture filter at any number of samples per symbol. However, using the RCP response implies that the bulk of the Nyquist channel characteristic resides in the demodulator, so matched filtering has been sacrificed for a simplified implementation that is effective in rejecting adjacent channels. Simulations have shown that this transmit/receive filter apportionment causes a degradation on the order of 0.5 dB in BER.

The frequency responses for the raised cosine pulse at 2, 3, 4, and 32 s/s are depicted in Figures 5a, b, c, and d, respectively. Observe that the ACR improves as the number of s/s is increased. Fortunately, at two s/s the analog anti-aliasing filter provides most of the needed ACR. Moreover, it is necessary to include additional integer sample rates in the demodulator between 2, 4, and 8 s/s, namely, 3 and 6 s/s to provide sufficient ACR. The relationship between sample and symbol rates as well as the number of s/s in the modulator and demodulator are listed in Tables 2a and 2b, respectively.

Table 2a. Modulator Rate Ranges
(*Msymbol/s, Msample/s*)

SYMBOL RATE	SAMPLES/SYMBOL	SAMPLE RATE
2.34375-4.6875	32	75-150
4.6875-9.375	16	75-150
9.375-18.75	8	75-150
18.75-37.5	4	75-150
37.5-75.0	2	75-150

Table 2b. Demodulator Rate Ranges
(*Msymbol/s, Msample/s*)

SYMBOL RATE	SAMPLES/SYMBOL	SAMPLE RATE
2.34375-4.6875	32	75-150
4.6875-6.25	24	112.5-150
6.25-9.375	16	100-150
9.375-12.5	12	112.5-150
12.5-18.75	8	100-150
18.75-25.0	6	112.5-150
25.0-37.5	4	100-150
37.5-50.0	3	112.5-150
50.0-75.0	2	100-150

To summarize, the pre-averager has several significant properties: 1) it serves as a variable rate FIR receive data filter of minimal complexity; 2) it reduces the processing rate and complexity of subsequent circuitry to 1 s/s; 3) it reduces the incoming noise bandwidth to approximately $\pm R_s/2$, thereby improving the input signal-to-noise (S/N) ratio established by the fixed analog anti-aliasing filter.

Data Detection

Data detection for the various modulation techniques is achieved with a memory table lookup of the even samples from the (I, Q) signal vector out of the pre-averagers. The sampling is synchronous and the symbol timing recovery loop will cause the even samples to automatically occur at the optimum data detection time instant. As stated previously, the largest memory size available at 75-MHz signaling speeds is 64K x 4, which provides for an I and Q input resolution of 8 bits.

Steady-State Recovery Loop Architecture


In 1977, a joint estimator-detector approach was developed at COMSAT Laboratories to provide an optimum way to recover carrier and clock for QPSK data transmission. It was found that the resultant technique which was dubbed Concurrent Carrier and Clock Synchronization (CCCS) applies to many types of digital data modulation. In particular, the CCCS technique is applicable to any modulation format that can be represented in quadrature carrier form: such as BPSK, QPSK, ... M-ary PSK, QAM, MSK, etc. Hence, this technique provides a basis for the PDM demodulator structure. Details of the CCCS technique are contained in References 2 and 3.

Some of the salient CCCS features which impact the PDM architecture will now be discussed. The CCCS method demonstrated that the optimum steady-state carrier phase and clock timing estimators are phase-locked loops (PLLs), which use post-detection feedback to remove data pattern noise and generate error signals that drive the loops. Post-detection data feedback is essentially noiseless because, even at a relatively poor BER of 10^{-2} , only 1 of every 100 detected data bits is incorrect. Hence, the loop S/N is merely reduced by a factor of 0.98 (-0.09 dB). Apart from knowing the transmitted data sequence, this is as well as a recovery loop can do.


For more complex signaling formats such as 8-, 16-PSK, and 16-QAM, where a quadrature carrier description of the IF signal requires several amplitude levels to be represented, the CCCS detected data feedback in the recovery loops must be multilevel. Multilevel feedback gives the larger average S/N samples proportionally more weight than the smaller ones, thereby maintaining the optimality of the recovery loop S/Ns. Moreover, the CCCS approach enables a common carrier, clock, and gain control recovery loop architecture to be used for any modulation format that can be represented in quadrature carrier form.

The basic error signal mechanism and loop filter for tracking in the CCCS architecture is illustrated in Figure 6. Table 3 lists the feedback signals needed for automatic gain control (AGC), carrier, and clock tracking. This common structure can be reconfigured in a MAC format by performing the multiplications sequentially and summing their products. Although this doubles the maximum speed requirement from 75 to 150 Msample/s, it is consistent with the speed already necessary for the pre-averager.


The error signals that drive the tracking loops are each processed by a loop filter to provide an output estimate.




a. 2 Samples/Symbol



b. 3 Samples/Symbol



c. 4 Samples/Symbol



d. 32 Samples/Symbol

Figure 5. Raised Cosine Pulse Filter Frequency Response

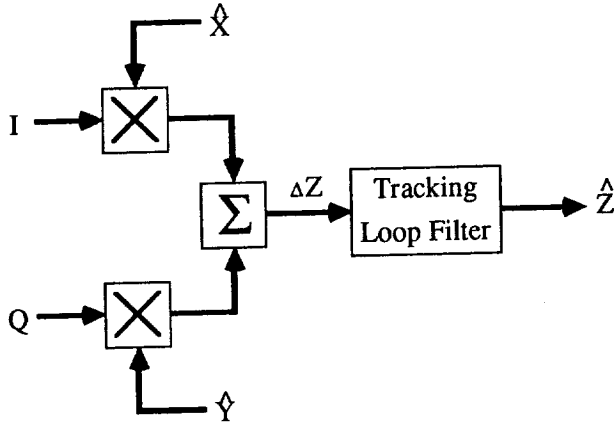


Figure 6. Recovery Loop Processor

Previous experience has shown that the AGC and clock loops need only be first order, whereas to track frequency offsets, the carrier loop must be second order. Hence, MACs will also be employed to satisfy the loop filter requirements. The loop bandwidth parameters can then be programmed by changing the multiplier gain constant. Moreover, the MAC architecture can be applied at numerous locations in the demodulator, including the acquisition circuitry.

Table 3. Tracking Loop Error Feedback Signals

FUNCTION	\hat{X}	\hat{Y}	\hat{Z}
Amplitude Level	\hat{I}	\hat{Q}	$\Delta \hat{A}$
Carrier Phase	\hat{Q}	$-\hat{I}$	$\hat{\theta}$
Symbol Timing	$\Delta \hat{I}$	$\Delta \hat{Q}$	$\hat{\tau}$

Notes: $\Delta \hat{A} = \hat{A} - A_{ref}$

$\Delta \hat{I} = \hat{I}(kT_s) - \hat{I}[(k-1)T_s]$

$\Delta \hat{Q} = \hat{Q}(kT_s) - \hat{Q}[(k-1)T_s]$

The output of the first order AGC loop filter is the estimate of the amplitude level error, $\Delta \hat{A}$, which is the control signal for the AGC amplifier. The AGC amplifier gain, G , is modeled as

$$G = \frac{G_{nom}}{1 + \Delta \hat{A}/A_{ref}}, \Delta \hat{A} > -A_{ref} \quad (4)$$

where G_{nom} is the nominal gain when $\Delta \hat{A} = 0$.

The output of the second order carrier loop filter is the estimate of the phase of the incoming signal. It includes the linear phase variations modulo 180° necessary to track carrier frequency offsets. Since a fixed frequency local oscillator (LO) is employed to down-convert the incoming signal to baseband, a carrier beat frequency occurs in the demodulated I and Q channels. A carrier phase rotator is used to eliminate the beat after the pre-averager data filters,

prior to detection. If the generalized incoming QAM signal is defined as

$$s[t, A, \theta(t), \tau] \triangleq A \{i(t, \tau) \cos[\omega t + \theta(t)] + q(t, \tau) \sin[\omega t + \theta(t)]\} \quad (5a)$$

where

A = incoming signal amplitude

ω = incoming signal frequency

$\theta(t)$ = incoming signal phase uncertainty

$i(t, \tau)$ = filtered in-phase modulating waveform

$q(t, \tau)$ = filtered quadrature modulating waveform

τ = modulating waveform timing uncertainty

and the quadrature LO outputs for down-conversions are

$$\text{lo}_I(t) = 2 \cos(\omega t) \quad (5b)$$

$$\text{lo}_Q(t) = 2 \sin(\omega t) \quad (5c)$$

The resulting baseband I and Q components prior to phase rotation are then

$$s_I(t) = A \{i(t, \tau) \cos[\theta(t)] + q(t, \tau) \sin[\theta(t)]\} \quad (6a)$$

$$s_Q(t) = A \{q(t, \tau) \cos[\theta(t)] - i(t, \tau) \sin[\theta(t)]\} \quad (6b)$$

To decouple the I and Q modulating waveforms, the carrier phase rotation is defined as

$$\begin{bmatrix} s'_I(t) \\ s'_Q(t) \end{bmatrix} = \begin{bmatrix} \cos[\hat{\theta}(t)] & -\sin[\hat{\theta}(t)] \\ \sin[\hat{\theta}(t)] & \cos[\hat{\theta}(t)] \end{bmatrix} \begin{bmatrix} s_I(t) \\ s_Q(t) \end{bmatrix} \quad (7a)$$

$$\begin{bmatrix} s'_I(t) \\ s'_Q(t) \end{bmatrix} = A \begin{bmatrix} i(t, \tau) \cos[\Delta\theta(t)] + q(t, \tau) \sin[\Delta\theta(t)] \\ q(t, \tau) \cos[\Delta\theta(t)] - i(t, \tau) \sin[\Delta\theta(t)] \end{bmatrix} \quad (7b)$$

where $\Delta\theta(t) = \theta(t) - \hat{\theta}(t)$, and the output estimate from the carrier loop filter is converted into two quadrature cosine and sine terms. The phase rotation described in equations (7) will also be implemented with MACs.

In the symbol timing tracking loop, the first order loop filter is actually a numerically controlled oscillator (NCO); which has an accumulator that holds the timing phase. Hence, the error signal from the timing phase detector is added with appropriate weighting to a constant that sets the nominal sample clock frequency, NR_s at the NCO input. The symbol clock as well as all other clocks used in the demodulator are then synchronously divided down from NR_s .

Burst-Mode Synchronization Techniques

To expedite lock and provide a high degree of false-and-miss detection reliability in burst mode, a parallel acquisition estimate path has been added to the tracking loop architecture. The initial carrier and clock phase as well as the amplitude level are estimated in this path and injected directly into the recovery loop accumulators. This effectively minimizes the loop lock-up transients. Since the accuracy of the acquisition measurement is proportional to the length of its observation interval, the burst false-and-miss detection probabilities can be made arbitrarily small.

In computing the acquisition estimates, it is desirable to uncouple them so they may be processed independently,

thereby having fewer degrees of uncertainty. For modulation techniques whose I and Q channels are not time staggered (such as offset formats), independent parallel processing of the estimates is possible with "01" modulation in both channels [4],[5]. The analog baseband I and Q signals defined in equations (6a and b) then may be described by

$$s_i(t) = \sqrt{2} A \sin[\pi R_s(t + \tau)] \{ \cos[\theta(t)] + \sin[\theta(t)] \} \quad (8a)$$

$$s_q(t) = \sqrt{2} A \sin[\pi R_s(t + \tau)] \{ \cos[\theta(t)] - \sin[\theta(t)] \} \quad (8b)$$

Equations (8a and b) can be reduced to

$$s_i(t) = 2A \sin[\pi R_s(t + \tau)] \sin[\theta(t) + \pi/4] \quad (9a)$$

$$s_q(t) = 2A \sin[\pi R_s(t + \tau)] \cos[\theta(t) + \pi/4] \quad (9b)$$

In the sampled domain, equations (9a and b) are rewritten as

$$I_{2k} \Delta 2A (-1)^k \cos(\phi_\tau/2) \sin(\theta + \pi/4) \quad (10a)$$

$$Q_{2k} \Delta 2A (-1)^k \cos(\phi_\tau/2) \cos(\theta + \pi/4) \quad (10b)$$

$$I_{2k-1} \Delta 2A (-1)^k \sin(\phi_\tau/2) \sin(\theta + \pi/4) \quad (10c)$$

$$Q_{2k-1} \Delta 2A (-1)^k \sin(\phi_\tau/2) \cos(\theta + \pi/4) \quad (10d)$$

where the timing phase offset, $\phi_\tau = 2\pi R_s \tau$, and the subscripts $2k$ and $2k-1$ denote even and odd samples, of the k th symbol, respectively.

Amplitude Level Acquisition Estimate

The most straightforward way to extract the amplitude A from equations (10a through d) independent of the phase and timing uncertainties is squaring, and then averaging to improve the estimate SNR. To simplify the hardware implementation and allow for sharing of common processing elements, the averaging should be done as soon as possible to lower the output sample rate. Because of the carrier frequency offset, the even and odd pairs of samples must be squared and combined in MACs on a symbol-by-symbol basis and then averaged.

$$E^2 \Delta \sum_k (I_{2k}^2 + Q_{2k}^2) = 4A^2 \cos^2(\phi_\tau/2) \quad (11a)$$

$$O^2 \Delta \sum_k (I_{2k-1}^2 + Q_{2k-1}^2) = 4A^2 \sin^2(\phi_\tau/2) \quad (11b)$$

Equations (11a and b) can then be combined to give the amplitude level estimate

$$\hat{A} = \sqrt{E^2 + O^2} / 2 \quad (12)$$

Equation 12 is most easily implemented as a memory table lookup. It was found in the emulations that 10 bits of resolution are needed for E^2 and O^2 because of the squaring. An intermediate compression table lookup is necessary to reduce the memory size in implementing equation (12) from 1 Mbyte to 64 kbytes.

Carrier Phase Acquisition Estimate

In reviewing equations (10a through d), it is apparent that there are several ways to isolate the carrier phase offset. For instance, the phase can be computed on a symbol-by-symbol basis as the arctangent of linear, square, or absolute value functions of I/Q , and then averaged; or I and Q can be squared first, and then averaged and processed as the arctangent of the sum of squares; or I and Q may be premultiplied by the preamble to remove the modulation, averaged, and the arctangent taken. All of these techniques have relative advantages and disadvantages. For instance, squaring the incoming samples increases the twofold ambiguity with "01" preamble modulation to fourfold; which either increases the complexity of the unique word detector or requires additional acquisition processing to unravel. Computing the arctangent on a symbol-by-symbol basis does not allow the arctangent processing element to be shared with the symbol timing loop. So the method chosen is the latter of the three examples for the following reasons. Premultiplication of the incoming samples by the known preamble removes the data modulation without S/N degradation. By next averaging the samples prior to the nonlinear arctangent operation, the S/N is improved. Finally, the largest pair of odd or even sample sums are chosen for the arctangent, so the twofold phase ambiguity is maintained. To make the odd vs even decision, the O^2 and E^2 sums, which were previously calculated in the amplitude level estimator are compared. Hence the resulting carrier phase estimate is computed from the ratio of I over Q samples as

$$\hat{\theta} = \tan^{-1} \left(\frac{\pm \sum |I_{2k} \text{ or } I_{2k-1}|}{\pm \sum |Q_{2k} \text{ or } Q_{2k-1}|} \right) - \pi/4 \quad (13)$$

where

$$E^2 \sum_{2k} > O^2 \sum_{2k-1}$$

Equation (13) will be implemented as a 64-kbyte memory table lookup.

To find the frequency offset, two such phase estimates are computed over the first and second halves of the preamble as θ_1 and θ_2 , respectively. The frequency offset can then be computed from the phase difference as

$$\Delta\omega = \frac{\Delta\theta}{\Delta T} = \frac{\theta_2 - \theta_1}{P/2} \quad (14)$$

where P is the total length of the preamble in symbol time units. The end-of-preamble phase estimate is determined from the measured phase and frequency difference as

$$\theta_{EOP} = \theta_2 + \Delta\omega \cdot \Delta T \quad (15)$$

Equations (14) and (15) will also be implemented as 64-kbyte memory table lookups.

Symbol Timing Acquisition Estimate

Again, there are several ways to compute the initial symbol timing error. It could be calculated from the arctangent of the square root of the previously computed values Q^2/E^2 , but the squaring would cause a twofold timing ambiguity which requires additional processing to resolve. It can also be computed from the arctangent of the largest pair of preamble premultiplied odd and even samples, which also requires an I^2 or Q^2 largest decision. The latter case turns out to be easier to implement since two of the tracking loop MACs are idle during acquisition and can be employed to calculate I^2 and Q^2 ; and in addition, the arctangent operation can be time shared with that required for carrier phase acquisition. So the symbol timing offset is computed from the ratio of odd over even samples as

$$\phi_\tau = 2 \tan^{-1} \left(\frac{\pm \sum |I_{2k-1} \text{ or } Q_{2k-1}|}{\pm \sum |I_{2k} \text{ or } Q_{2k}|} \right) \quad (16)$$

where

$$I^2 \Delta \sum_k (I_{2k}^2 + I_{2k-1}^2) \stackrel{I}{\geq} Q^2 \Delta \sum_k (Q_{2k}^2 + Q_{2k-1}^2)$$

Equation (16) will share the same 64-kbyte memory table as the carrier phase in equation (13). The slight differences in the expressions will be compensated for in the end of preamble phase computation from equations (14) and (15).

CONCLUSIONS

Operation of digital signal processing (DSP) circuitry at sample rates as high as 150 MHz appears feasible. The two most speed-critical areas are memories and multiplier-accumulators. Currently available high-density static RAMs can only operate up to approximately 80 MHz and must be ping-ponged to achieve the desired rate. The workhorse of the processing is clearly the multiplier-accumulator. To achieve 150-MHz operation with sufficient margin and power efficiency, GaAs is the most appropriate technology; potential GaAs vendors have recommended a standard-cell rather than a gate-array approach for this application.

Subsequent hardware emulations have verified the fundamental design approach presented in this paper, as well as the bit resolutions and aperture lengths used. The results will be submitted in a forthcoming publication.

ACKNOWLEDGMENTS

The author would like to acknowledge J. Thomas for his original contributions in this area from COMSAT's Digitally Implemented Modem Program, and F. Faris for developing the hardware emulation program.

REFERENCES

- [1] J. J. Poklemba, C. Wolejsza, and J. Thomas, "Programmable Noise Bandwidth Reduction by Means of Digital Averaging," COMSAT Invention Disclosure No. 25-E-14, June 27, 1986.
- [2] J. J. Poklemba, "Concurrent Carrier and Clock Synchronization for Data Transmission Systems," U.S. Patent No. 4, 419, 759, December 6, 1983.
- [3] J. J. Poklemba, "A Joint Estimator-Detector for QPSK Data Transmission," COMSAT Technical Review, Vol. 14, No. 2, Fall 1984, pp. 211-259.
- [4] S. A. Rhodes, "Synchronization for TDMA/QPSK Up-Link Transmissions With Digital On-Board Processing," COMSAT Technical Memorandum, CL-15-86, August 13, 1986.
- [5] S. A. Rhodes, "Modification of Acquisition Algorithms for On-Board Demodulator," COMSAT Technical Note, CTD-87/007, January 14, 1987.

